Jess: Making Your Own Rules with Sandia's Rule Engine

Many problems are naturally algorithmic, meaning that a well-defined series of steps leads to a solution. Most computer applications fall into this category. However, some problems, such as the following examples, resist being reduced to rote computation: "Is our network under attack?" "Is this document fraudulent?" "How should we schedule our resources?" These problems have heuristic solutions, best described by expert knowledge and "rules of thumb," but such solutions are impossible to reduce to tractable equations. When you are faced with this type of problem, sometimes you need to make your own rules—using Sandia's Jess rule engine, of course.

Written in Sun Microsystems' Java programming language, Jess is a portable library for symbolic computation. Rule-based systems written with Jess have the ability to rapidly process large amounts of symbolic or numeric data while providing emergent solutions to ill-conditioned, nonalgorithmic problems. Such systems can stand in for the human element in a larger simulation, or they can provide decision support based on live data fusion in real situations as they unfold.

Sandia scientist Ernest Friedman-Hill originally developed Jess to be used as the "brain" of a mobile agent—a smart software program that travels around a network as it runs, bringing its internal state with it. The agent was designed to help solve an information-security problem known as mosaic classification, or the possibility of creating classified information by combining multiple pieces of unclassified information. The solution involved monitoring the computer-usage patterns of individual users.

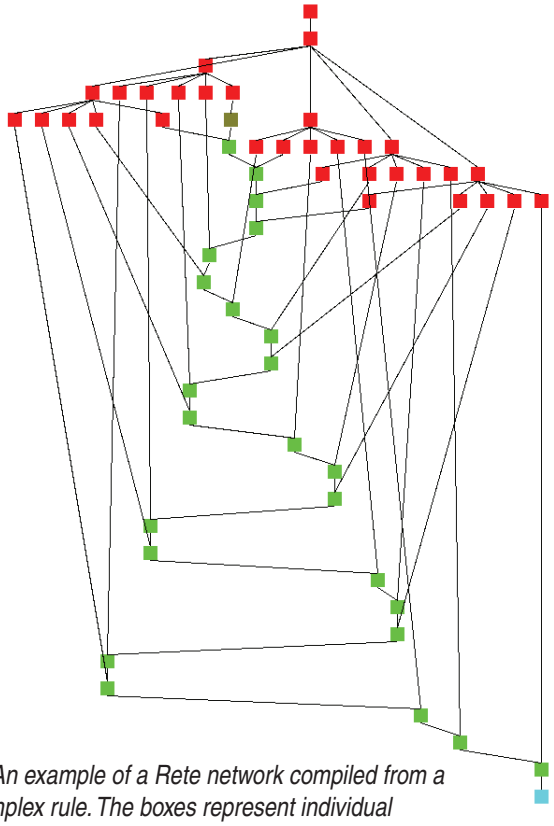

Figure 1. An example of a Rete network compiled from a single complex rule. The boxes represent individual discriminations, and the connections represent data paths.

To write a program with Jess, you specify heuristic knowledge in the form of rules. Jess then applies these rules to the data at hand. Jess rules can be expressed either in a special rule language based on Lisp, or in XML. Jess's rule language offers full access to all the capabilities of the Java language and its libraries.

SAND2006-2560P
April 2006

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Jess uses an advanced version of the Rete algorithm [1] to build a discriminator network from the user's rules. This network can process millions of data items in fractions of a second. The innovations making this possible include a novel indexing scheme that accelerates correlations between data items in the network by filtering out irrelevant data.

A rule engine, like a relational database, processes large amounts of information in the form of tables; each table consists of multiple columns. A database developer can speed up database queries by adding custom-designed indices, which let the database engine find information quickly. These indices often use compound keys, which consist of multiple data items from multiple columns and tables. Jess creates custom indices automatically, based on the user's rules, making operation simpler and data access faster.

What began as a minor component of an information-security research project has become one of Sandia, California's most successful intellectual-property licensing properties. Jess has been licensed to hundreds of commercial concerns, including dozens of Fortune 500 companies. Jess-based systems are currently deployed in a number of industries, including finance, computer network security, transportation scheduling, inventory control, and manufacturing process flow.

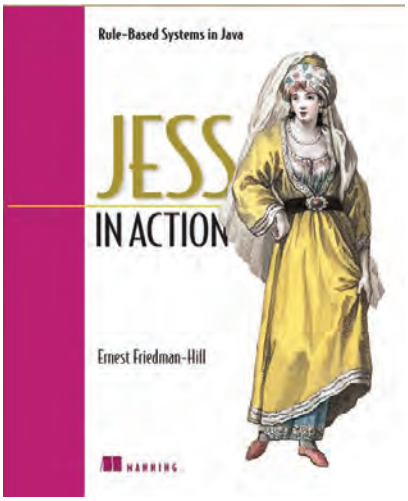

Figure 2. Sandia scientist Ernest Friedman-Hill wrote the book on Jess.

One of the most sophisticated commercial applications of Jess is in a freight-car routing application deployed by a major U.S. railroad. The application is a management information decision support system. It tracks thousands of freight cars, both those owned by the railroad and those borrowed from or lent to other railroads, in real time. The application tracks the cars' physical characteristics (capacity, size, and capabilities like refrigeration) and their location, destination, history, service time, and other qualities. The application then makes recommendations about how to route cars from place to place so that materials arrive in the right place on time and so that borrowed and lent cars return to their proper location. At the same time, the application keeps track of locomotive capacity and other factors that could affect the railroad's routine. This enormous system continuously processes a huge amount of data and provides a mission-critical service for a substantial company.

Sandia has also granted thousands of academic licenses for Jess. Jess is used in many university courses on artificial intelligence all over the world.

Jess has evolved considerably over the years. The modern version of Jess includes an integrated development environment (IDE) based on the Eclipse [2] platform. The IDE includes an editor, a debugger, and visualization tools.

Current work involves a collaboration with the University of Alabama at Birmingham to bring Jess to a handheld platform equipped with a global positioning system and to develop an application to guide first responders in emergent situations.

You can learn more about Jess at http://www.jessrules.com.

[1] Forgy, Charles L. "Rete: A fast algorithm for the many pattern/many object pattern match problem." *Artificial Intelligence*, 1982, 19:17–37.

[2] Eclipse Web site, http://www.eclipse.org.

For more information, contact
Ernest Friedman-Hill, ejfried@sandia.gov